

Performance Validation Test

Cisco MDS 9509

By



at

Spirent SmartLabs, Calabasas California - December 2002

for

Cisco Systems, Incorporated

Miercom
www.mier.com
410 Hightstown Road
Princeton Junction., New Jersey 08550
609-490-0200 fax 609-490-0610

Michael B Hommer, Senior Manager of Lab Testing
Robert J Smithers, President

Table of Contents

Executive Summary	3
Purpose.....	3
Configuration	3
Summary of Findings.....	3
Fully Configured Throughput and Latency Test.....	4
Many-to-One Throughput Test	6
Many-to-Few Fan-In Test	8
Head-of-Line Blocking Test	11
Enhanced Head-of-Line Blocking Test.....	13
Cold Reboot Test	15
Non-Disruptive Code Load and Activation Test.....	16
Load Balancing on Redundant ISLs Test.....	17
Redundant ISL Failover Test	19
ISL Fairness Test	20
32-Port Line Card Oversubscription and Fairness Test	22

Table of Tables

Table 1: Fully Configured, Throughput and Latency.....	5
Table 2: Fully Configured, Latency Distribution	5
Table 3: Many-to-One, Throughput	6
Table 4: Many-to-One, Latency Distribution.....	7
Table 5: Many-to-One, Throughput Distribution.....	7
Table 6: Many-to-Few Fan-In, Throughput	9
Table 7: Many-to-Few Fan-In, Throughput Distribution Across Two Cards	10
Table 8: Many-to-Few Fan-In, Throughput Distribution Across Four Cards	10
Table 9: Head-of-Line Blocking, Throughput and Latency	11
Table 10: Head-of-Line Blocking, Throughput and Latency w/ Different Frame Sizes	11
Table 11: Head-of-Line Blocking, Latency Distribution w/ 60-Byte Frames.....	12
Table 12: Head-of-Line Blocking, Latency Distribution w/ 2148-Byte Frames.....	12
Table 13: Enhanced Head-of-Line Blocking, Throughput and Latency	13
Table 14: Enhanced Head-of-Line Blocking, Latency Distribution w/ 60-Byte Frames	14
Table 15: Enhanced Head-of-Line Blocking, Latency Distribution w/ 2148-Byte Frames	14
Table 16: Cold Reboot	15
Table 17: Non-Disruptive Code Load and Activation	16
Table 18: Load Balancing on Redundant ISLs	18
Table 19: Redundant ISL Failover	19
Table 20: ISL Fairness, Throughput and Latency.....	20
Table 21: ISL Fairness, Latency Distribution	21
Table 22: 32-Port Line Card Oversubscription and Fairness	23

Executive Summary

Purpose

Miercom was contracted by Cisco Systems, Incorporated (Cisco) to test and observe their MDS 9509 storage area network (SAN) director-class Fibre Channel switch. The purpose was to test the overall characteristics and performance of the product. Testing was conducted at the Spirent Communications facilities in Calabasas, California during the week of December 16, 2002.

Configuration

Most testing was conducted using one or two MDS 9509 directors. The MDS 9509 directors were configured with two redundant supervisor/fabric cards unless otherwise noted. Fibre Channel Class 3 traffic was passed through each director at two gigabits per second (Gb/s). At most, 112 Fibre Channel ports were used during testing with unidirectional or bi-directional traffic flows. Throughput, frame loss, latency and latency distribution statistics were collected for most tests. Two MDS 9216 Fabric switches, one in a 32-port configuration and one in a 48-port configuration, were used in some tests to simulate a core-edge model.

Summary of Findings

The performance on the MDS 9509 was very good. For all tests using large frames, full line rate was achieved. Using bi-directional flows with small frames, there was some degradation from 100% line rate. However, the throughput was still much higher than had been previously observed in other directors.

The failover capabilities and non-disruptive code load functionality of the MDS 9509 were near the top. The reboot time was just over two minutes and the supervisor failover was instantaneous. The code load capabilities for a single supervisor also proved non-disruptive and took less than 15 seconds to complete.

Overall the MDS 9509 proved to be a non-blocking, fully redundant architecture with excellent performance in every topology tested.

Michael Hommer
Manager of Consulting Services

Fully Configured Throughput and Latency Test

Objective: Determine the maximum throughput of the Device Under Test (DUT) in a one-to-one, initiator-to-target scenario.

Test Bed:




Figure 1: Fully Configured Throughput and Latency Test

Five Spirent SmartBits 6000 chasses were connected to the DUT on 112 2-Gb/s Fibre Channel ports. Bi-directional traffic was sent through 56 pairs of ports in the DUT to yield 112 unique flows. Latency tests ran for 10 seconds and the throughput tests ran for 30 seconds using both 60- and 2148-byte frames, at 10% and 100% intended loads. SmartFabric gathered throughput, frame loss, latency and latency distribution statistics.

Results:

Table 1: Fully Configured, Throughput and Latency

Frame Size (Bytes)	Intended Load (%)	Latency		Throughput (MB/s)		Load Delta (%) ¹	Actual Load (%)	Lost Frames
		Min. (μs)	Max. (μs)	Per Port	Total			
60 1 sup	10	7.1	8.0					
	100	7.2	52.9	149.77	16774.56	1.33	98.67	0
2148 1 sup	10	19.6	28.9					
	100	19.7	218.9	210.15	23536.99	0.00	100.00	0
2148 2 sup	10	19.6	27.2					
	100	19.7	230.4					

Table 2: Fully Configured, Latency Distribution

Frame (bytes)	Intended Load (%)	Bucket Sizes (Number of Frames)						
		≤ 10 μs	≤ 20 μs	≤ 30 μs	≤ 40 μs	≤ 50 μs	≤ 100 μs	≤ 500 μs
60 1 sup	10	283,333,120	0	0	0	0	0	0
	100	1,886	234,656,698	127,503,028	2,408,986,494	62,177,325	6,889	0
2148 1 sup	10	0	8,672,597	2,284,363	0	0	0	0
	100	0	2,309,643	58,608,668	13,481,826	10,450,351	20,897,232	3,828,600
2148 2 sup	10	0	9,057,496	1,899,464	0	0	0	0
	100	0	5,930,815	56,681,508	13,643,963	9,356,407	20,515,335	3,448,292

Conclusion:

The MDS 9509 showed excellent performance regardless of frame size. It achieved near line rate with small frames (98.67%) and full line rate with large frames, both with 100% intended load. Regardless of the load, the minimum latency for both frame sizes was very consistent. For small frames, it varied from 7.2 to 52.9 μs under 100% intended load. For large frames, the latency ranged from 19.7 to 218.9 μs under 100% intended load. Furthermore, the MDS 9509 was able to sustain traffic at a much higher throughput rate for minimum- and maximum-sized frames while maintaining a more consistent latency than other switches tested.

When a second supervisor/fabric card was added, there was no impact on the overall results. While the latency distribution is fairly close between using one or two supervisors/fabrics, there are slightly more frames in the sub-20 μs range when using two supervisor/fabrics.

¹ Load Delta is a term used by the test tool application SmartFabric. It is defined as the difference between the configured or intended load and the actual load received.

Many-to-One Throughput Test

Objective: Determine the throughput of a DUT with traffic from 15 initiators to 1 target.

Test bed:




Figure 2: Many-to-One Throughput Test

Unidirectional traffic was generated by the Spirent SmartBits 6000 into the DUT on fifteen 2-Gb/s Fibre Channel ports, and directed to the same Spirent SmartBits 6000 via one 2-Gb/s Fibre Channel port. The throughput tests ran for 30 seconds and the latency tests ran for 10 seconds, using 60- and 2148-byte frames with an intended load of 100%. SmartFabric gathered throughput, frame loss, latency and latency distribution statistics. All ports were on a single blade of the DUT.

Given the 15-to-1 ratio of ingress to egress ports, with 100% intended load the ideal load per ingress port would be approximately 6.66%

Results:

Table 3: Many-to-One, Throughput

Frame Size (Bytes)	Throughput (MB/s)	Intended Load (%)	Load Delta (%)	Ideal ² Load (%)	Actual Load (%)	Latency	
						Min. (µs)	Max. (µs)
60	152.02	100	93.32	6.66	6.68	7.4	542.1
2148	210.48	100	93.32	6.66	6.68	19.8	5255.8

² Ideal load is based on a calculation using a perfect Fibre Channel output clock frequency. The actual frequency may be higher or lower within +/- 100 PPM per output port. Because of this tolerance, the actual load may be higher than the ideal load.

Table 4: Many-to-One, Latency Distribution

Frame Size (bytes)	Bucket Sizes (Number of frames)						
	$\leq 10 \mu s$	$\leq 50 \mu s$	$\leq 100 \mu s$	$\leq 300 \mu s$	$\leq 400 \mu s$	$\leq 500 \mu s$	$\leq 550 \mu s$
60	6	28,199	27,152	256,889	29,085,603	12,964	4,836,628
	$\leq 10 \mu s$	$\leq 50 \mu s$	$\leq 100 \mu s$	$\leq 500 \mu s$	$\leq 1000 \mu s$	$\leq 5000 \mu s$	$\leq 5500 \mu s$
2148	0	3	5	42	2,256	13,600	1,307,808

Table 5: Many-to-One, Throughput Distribution

Initiator Port	Throughput (MB/s)	
	60-Byte Frames	2148-Byte Frames
1/1	10.13	14.03
1/2	10.13	14.03
1/3	10.13	14.03
1/4	10.13	14.03
1/5	10.13	14.03
1/6	10.13	14.03
1/7	10.13	14.03
1/8	10.13	14.03
1/9	10.13	14.03
1/10	10.13	14.03
1/11	10.14	14.03
1/12	10.14	14.03
1/13	10.14	14.04
1/14	10.14	14.04
1/15	10.15	14.04

Conclusion:

The MDS 9509 exhibited excellent throughput and latency for all frame sizes in this test. Both large and small frames achieved the same percentage of throughput with latencies under 550 μs for small frames and under 5,300 μs for large frames. More impressive was the distribution of the traffic flows, which was varied +/- 0.01 MB/s for small frames and +/- 0.005 MB/s for large frames.

Many-to-Few Fan-In Test

Objective: Test for architectural blocking between line components within a DUT.

Test Bed:




Figure 3: Many-to-Few Fan-In Test Across Two Cards

Many-to-Few Fan-In Test Across Two Cards

Four groups of four 2-Gb/s Fibre Channel ports were connected from one Spirent SmartBits 6000 into 16 ports on a single line card on the DUT, identified as Set 1. Four 2-Gb/s Fibre Channel ports from another Spirent SmartBits 6000 were then connected to a second line card on the DUT. This was known as Set 2. A flow was initiated from the first port of each group in Set 1 to the first port in Set 2. Another flow was established from the second port of each group in Set 1 to the second port in Set 2. An additional flow was started from the third port of each group in Set 1 to the third port in Set 2. Lastly, a flow was generated from the fourth port of each group in Set 1 to the fourth port in Set 2. All flows were unidirectional, used 60- and 2148-byte frames, lasted 10 seconds, and were configured for an intended load of 100%. SmartFabric gathered throughput statistics.




Figure 4: Many-to-Few Fan-In Test Across Four Cards

Many-to-Few Fan-In Test Across Four Cards

A second topology was also tested. The ports for Set 1 were spread across the first four line cards in groups of four. The ports for Set 2 were connected to the first line card. Flows were established in the same manner described above in the many-to-few fan-in test across two cards.

Results:

Table 6: Many-to-Few Fan-In, Throughput

Frame Size (bytes)	Ideal ³ Throughput (MB/s)	Actual Throughput (MB/s)	
		Across 2 Cards	Across 4 Cards
60	607.14	607.92	608.03
2148	840.61	841.89	841.48

³ Ideal throughput is based on a calculation using a perfect Fibre Channel output clock frequency. The actual frequency may be higher or lower within +/- 100 PPM per output port. Because of this tolerance, the actual throughput may be higher than the ideal throughput.

Table 7: Many-to-Few Fan-In, Throughput Distribution Across Two Cards

Initiator Port	Throughput (MB/s)	
	60-Byte Frames	2148-Byte Frames
1/1	37.97	52.58
1/2	37.99	52.61
1/3	38.03	52.67
1/4	37.97	52.58
1/5	37.97	52.59
1/6	37.99	52.61
1/7	38.04	52.68
1/8	37.98	52.59
1/9	37.97	52.59
1/10	37.99	52.61
1/11	38.04	52.68
1/12	37.98	52.59
1/13	37.98	52.60
1/14	38.00	52.62
1/15	38.04	52.68
1/16	37.98	52.60
Average	38.00	52.62
Ideal	37.95	52.54
Range	+/- 0.035	+/- 0.045
Total	607.92	841.89

Table 8: Many-to-Few Fan-In, Throughput Distribution Across Four Cards

Initiator Port	Throughput (MB/s)	
	60-Byte Frames	2148-Byte Frames
1/1	37.97	52.60
1/2	37.97	52.60
1/3	38.02	52.57
1/4	38.03	52.57
2/1	37.97	52.60
2/2	37.97	52.60
2/3	38.03	52.58
2/4	38.03	52.58
3/1	37.97	52.61
3/2	37.97	52.61
3/3	38.03	52.58
3/4	38.03	52.58
4/1	37.98	52.61
4/2	37.98	52.61
4/3	38.03	52.59
4/4	38.04	52.59
Average	38.00	52.59
Ideal	37.95	52.54
Range	+/- 0.035	+/- 0.02
Total	608.03	841.48

Conclusion:

Regardless of the topology tested, results were almost identical. For small frames in both configurations, the average throughput and variance were identical. For large frames there was an insignificant difference between the two topologies. Overall the distribution of traffic was even between all flows.

Head-of-Line Blocking Test

Objective: Test the impact on an uncongested port caused by congestion on an uncongested port. One of the input ports is configured to send traffic to both the congested port and the uncongested port and the throughput impact is measured.

Test Bed:




Figure 5: Head-of-Line Blocking Test

Four 2-Gb/s Fibre Channel ports were connected from one Spirent SmartBits 6000 into one group of 4 ports of the DUT. The first port sent traffic to the second port (Flow A) while the third port generated traffic to both the second port (Flow B) and fourth port (Flow C). All traffic is unidirectional. Throughput tests ran for 30 seconds and the latency tests ran for 10 seconds, all at 100% intended load using 60- and 2148-byte frames. SmartFabric gathered throughput, frame loss, latency and latency distribution statistics.

Results:

Table 9: Head-of-Line Blocking, Throughput and Latency

Flow	Intended Load (%)	Ideal Load (%)	Actual Load (%)	60-Byte Frames			2148-Byte Frames		
				Throughput (MB/s)	Latency		Throughput (MB/s)	Latency	
					Min. (µs)	Max. (µs)		Min. (µs)	Max. (µs)
A	100	50	50	79.46	7.7	105.5	105.14	19.8	870.8
B	50	50	50	72.36	7.3	95.6	105.08	30.0	231.0
C	50	50	50	72.36	7.4	10.8	105.08	19.7	21.1
Total		150	150	224.18			315.30		

Table 10: Head-of-Line Blocking, Throughput and Latency w/ Different Frame Sizes

Flow	Frame Size (Bytes)	Intended Load (%)	Ideal Load (%)	Actual Load (%)	Throughput (MB/s)
A	2148	100	50	50	107.93
B	60	50	50	50	73.86
C	60	50	50	50	73.86
Total			150	150	255.65

Table 11: Head-of-Line Blocking, Latency Distribution w/ 60-Byte Frames

Flow	Bucket Sizes (Number of Frames)						
	$\leq 10 \mu s$	$\leq 20 \mu s$	$\leq 30 \mu s$	$\leq 40 \mu s$	$\leq 50 \mu s$	$\leq 100 \mu s$	$\leq 110 \mu s$
A	6	51	6,330,417	947,501	706	11,906,151	1,984,072
B	6	24	27	24	27	12,648,697	0
C	12,578,100	70,705	0	0	0	0	0

Table 12: Head-of-Line Blocking, Latency Distribution w/ 2148-Byte Frames

Flow	Bucket Sizes (Number of Frames)						
	$\leq 10 \mu s$	$\leq 20 \mu s$	$\leq 50 \mu s$	$\leq 100 \mu s$	$\leq 250 \mu s$	$\leq 500 \mu s$	$\leq 1000 \mu s$
A	0	1	4	9	20	32	818,777
B	0	0	2	6	489,172	0	0
C	0	447,589	41,591	0	0	0	0

Conclusion:

For both large and small frames, the MDS 9509 did not hinder traffic destined for the uncongested port. The expected traffic performance was achieved for all flows. Of particular note is the latency of Flow C which remains low despite the second half of the stream, Flow B, going to a congested port.

An additional test was run using mixed frames sizes to demonstrate that similar results are achieved even when the streams are not identical as shown in table 10.

Enhanced Head-of-Line Blocking Test

Objective: Test the impact on an uncongested port caused by congestion on an uncongested port. One of the input ports is configured to send traffic to both the congested port and the uncongested port and the throughput impact is measured. In addition, a third input port is used to add further congestion to the congested port.

Test Bed:




Figure 6: Enhanced Head-of-Line Blocking Test

Four 2-Gb/s Fibre Channel ports were connected to one Spirent SmartBits 6000 into one group of four contiguous ports on the DUT. An additional 2-Gb/s SmartBits port was connected to another port of the DUT. The first port sent traffic to the second port (flow A), while the third port generated traffic to both the second (flow B) and fourth port (flow C). In addition, the fifth port also sent traffic to the second port (flow D). The throughput tests ran for 30 seconds and the latency tests ran for 10 seconds, all at 100% intended load, using 60-byte frames and 2148-byte frames. SmartFabric gathered throughput, frame loss, latency and latency distribution statistics.

Results:

Table 13: Enhanced Head-of-Line Blocking, Throughput and Latency

Flow	Intended Load (%)	Ideal Load (%)	Actual Load (%)	60-Byte Frames			2148-Byte Frames		
				Throughput (MB/s)	Latency		Throughput (MB/s)	Latency	
					Min. (µs)	Max. (µs)		Min. (µs)	Max. (µs)
A	100	33.3	33.3	50.61	7.4	143.5	70.06	30.0	1208.1
B	50	33.3	33.3	50.61	8.2	132.0	70.06	19.8	1193.8
C	50	33.3	33.3	50.61	7.1	10.8	70.06	19.7	21.0
D	100	33.3	33.3	50.61	7.8	143.5	70.06	40.2	1208.0
Total		133.3	133.3	202.44			280.24		

Table 14: Enhanced Head-of-Line Blocking, Latency Distribution w/ 60-Byte Frames

Flow	Bucket Sizes (Number of Frames)						
	$\leq 10 \mu s$	$\leq 20 \mu s$	$\leq 30 \mu s$	$\leq 40 \mu s$	$\leq 50 \mu s$	$\leq 100 \mu s$	$\leq 150 \mu s$
A	1	17	14	26	0	84	11,272,049
B	3	8	8	7	9	27,689	11,272,052
C	11,145,482	154,293	0	0	0	0	0
D	6	23	14	56,694	671	27,858	11,276,416

Table 15: Enhanced Head-of-Line Blocking, Latency Distribution w/ 2148-Byte Frames

Flow	Bucket Sizes (Number of Frames)						
	$\leq 20 \mu s$	$\leq 50 \mu s$	$\leq 100 \mu s$	$\leq 500 \mu s$	$\leq 1000 \mu s$	$\leq 1200 \mu s$	$\leq 1300 \mu s$
A	1	1	3	20	24	20	435,586
B	0	1	2	39	1,135	435,580	0
C	435,661	1,095	0	0	0	0	0
D	0	1	3	20	3,332	40	435,587

Conclusion:

The Cisco MDS 9509 evenly distributed the available bandwidth to all flows. The congestion did not affect the latency of the uncongested port which was shown as Flow C in tables 14 and 15.

Cold Reboot Test

Objective: Observe the DUT's resilience to power outages.

Test Bed:




Figure 7: Cold Reboot Test

Two 2-Gb/s SmartBits Fibre Channel ports were attached to the DUT. Using the Fibre Channel Group Control function within SmartWindow, the status of the ports was monitored. Power was interrupted for 10 seconds and the timestamp in the SmartWindow Group Control window was noted when power was restored. When the ports reached the “Ready to Test” status, the timestamp was recorded again and a test initiated to confirm that the switch was operational.

Results:

Table 16: Cold Reboot

Trial	Total Down Time (min:sec)
1	2:05
2	2:05
3	2:05

Conclusion:

During a reboot of the switch, the total down time was approximately two minutes. All fabric services and other operations resumed operation without intervention.

Non-Disruptive Code Load and Activation Test

Objective: Verify the ability of the DUT to have software upgraded without impact.

Test Bed:




Figure 8: Non-Disruptive Code Load and Activation Test

Five Spirent SmartBits 6000 chasses were connected into the DUT on 112 2-Gb/s Fibre Channel ports. Unidirectional traffic was sent through 56 pairs of ports in the DUT to create 112 unique flows. These flows ran for 60 seconds with 60-byte frames at 100% offered load. In the middle of the test, a code load was initiated and completed. SmartFabric gathered throughput and frame loss statistics to determine if any throughput impairments occurred during the code load. The load time was based on the time from which the code load command was issued to the time that the command prompt returned.

Results:

Table 17: Non-Disruptive Code Load and Activation

Trial	Number of Supervisors	Total Impact (μs)	Total Time to Upload Code (min:sec)
1	1	0.0	Less than 0:15
2	2	0.0	Less than 0:15

Conclusion:

Code upgrades and downgrades were successful with one and two supervisor cards present. In both cases there was no disruption to the traffic and the code load was completed in 15 seconds.

The ability to run one supervisor with a particular software version and the other supervisor with a different software version was demonstrated with no impact to switching functionality. This feature is particularly useful as it allows a software upgrade to be performed and verified on the standby supervisor and reversed, if desired, without any impact to the running version of software.

Load Balancing on Redundant ISLs Test

Objective: Determine the efficacy of redundant ISLs and the load-balancing capability of the MDS 9500 director using the PortChannel feature.

Test Bed:




Figure 9: Load Balancing on Redundant ISLs Test

Two 2-Gb/s Fibre Channel inter-switch links (ISLs) were connected between two DUTs, an MDS 9509 and an MDS 9216, using the PortChannel feature. Six 2-Gb/s Fibre Channel ports were connected from the Smartbits to DUT A. Two 2-Gb/s Fibre Channel ports were connected from the Smartbits to DUT B. Bi-directional flows were started between the six ports on DUT A to the two ports on DUT B. The throughput tests ran for 30 seconds using 60-byte frames and then 2148-byte frames. SmartFabric gathered throughput statistics.

Results:

Table 18: Load Balancing on Redundant ISLs

Flow	Throughput (MB/s)	
	60-Byte Frames	2148-Byte Frames
A1 → B1	23.08	34.90
A1 → B2	23.08	34.90
A2 → B1	23.08	34.90
A2 → B2	23.08	34.90
A3 → B1	23.08	34.90
A3 → B2	23.08	34.90
A4 → B1	23.08	34.90
A4 → B2	23.08	34.90
A5 → B1	23.08	34.90
A5 → B2	23.08	34.90
A6 → B1	23.08	34.90
A6 → B2	23.08	34.90
B1 → A1	23.07	34.90
B1 → A2	23.07	34.90
B1 → A3	23.07	34.90
B1 → A4	23.07	34.90
B1 → A5	23.07	34.90
B1 → A6	23.07	34.90
B2 → A1	23.08	34.90
B2 → A2	23.08	34.90
B2 → A3	23.08	34.90
B2 → A4	23.08	34.90
B2 → A5	23.08	34.90
B2 → A6	23.08	34.90
Total	553.89	837.62

Conclusion:

For both frame sizes, the redundant ISL bandwidth was evenly distributed to all flows.

Redundant ISL Failover Test

Objective: Determine the failover abilities of redundant, load balanced ISL links using the PortChannel Feature.

Test Bed:




Figure 10: Redundant ISL Failover Test

Two 2-Gb/s Fibre Channel inter-switch links (ISLs) were connected between the two DUTs (MDS 9509 and an MDS 9216) and configured as a PortChannel. A 2-Gb/s Fibre Channel port was connected from the SmartBits to each DUT. The SmartFabric failover test was run for 30 seconds using 2148-byte frames at 100% intended load. The failover test statistics were collected for each run.

Results:

Table 19: Redundant ISL Failover

Trial	Downtime (ms)
1	185.81
2	121.71
3	125.43
4	144.86
Average	144.45

Conclusion:

Service interruptions to traffic across ISL links averaged 144 milliseconds. Traffic flowing across the non-disrupted link is not affected by the failover. The removal of either link did not cause an FSPF routing update to disrupt traffic.

ISL Fairness Test

Objective: Determine the DUT's ability to distribute ISL bandwidth evenly to ingress ports.

Test Bed:




Figure 11: ISL Fairness Test

One 2-Gb/s Fibre Channel inter-switch link (ISL) was connected between the two DUTs. Two 2-Gb/s Fibre Channel ports were connected to each DUT from the Smartbits 6000 as port pairs, port A to port B and port C to port D. Using SmartFabric, throughput tests were configured with bi-directional flows between the port pairs. The latency tests were configured using unidirectional flows for the ports pairs. All tests were run at 100% offered load using 60-byte frames and then 2148-byte frames.

Results:

Table 20: ISL Fairness, Throughput and Latency

Frame Size (Bytes)	Flow	Throughput (MB/s)	Latency	
			Min. (µs)	Max. (µs)
60	1 (A → B)	75.56	14.2	107.8
	2 (B → A)	75.56		
	3 (C → D)	75.57	14.6	108.8
	4 (D → C)	75.57		
2148	1 (A → B)	105.09	48.2	880.0
	2 (B → A)	105.09		
	3 (C → D)	105.09	38.1	881.1
	4 (D → C)	105.09		

Table 21: ISL Fairness, Latency Distribution

Frame Size (Bytes)	Bucket Sizes (Number of Frames)						
	≤ 20 μs	≤ 30 μs	≤ 40 μs	≤ 70 μs	≤ 90 μs	≤ 100 μs	≤ 110 μs
60	15	45	58,712	1,897	192	28,996,799	4,832,791
	≤ 30 μs	≤ 500 μs	≤ 600 μs	≤ 700 μs	≤ 800 μs	≤ 850 μs	≤ 900 μs
2148	0	91	2,300	30	30	15	1,307,876

Conclusion:

The MDS 9509 very evenly distributed all traffic attempting to traverse the ISL. The latency is also fairly balanced across the flows.

32-Port Line Card Oversubscription and Fairness Test

Objective: Measure the oversubscription ratio and verify non-blocking performance of the MDS 9000 Family 32-port line card.

Test Bed:




Figure 12: 32-Port Line Card Oversubscription and Fairness Test

The ports on the 32-port module were grouped into eight groups of four ports each, A through H. The ports on the 16-port module were grouped into eight groups of two ports each. Each of the ports in a group on a 32-port line card transmitted to two ports in a group on a 16-port line card as shown in the figure. A total of 48 Fibre Channel ports on the Smartbits 6000 were used. A SmartFabric throughput test was configured for 10 seconds at 100% intended load using 2148-byte frames. SmartFabric gathered the throughput statistics.

Results:

Table 22: 32-Port Line Card Oversubscription and Fairness Test

Group	Throughput (MB/s)	Actual Load (%)
A	248.60	29.57
B	248.64	29.58
C	248.68	29.58
D	248.75	29.59
E	248.76	29.59
F	248.80	29.60
G	248.58	29.57
H	248.62	29.58

Conclusion:

The oversubscription ratio of each group of four ports obtained through this test was 3.3 to 1. The test proved that the 32-port module is non-blocking because throughput was evenly distributed across all port groups and there were no dropped frames.

Glossary

Backpressure

Any technique used by a DUT to attempt to avoid frame loss by impeding external sources of traffic from transmitting frames to congested interfaces (RFC 2285). This term was originally used to describe a primitive form of congestion control used in half-duplex Ethernet LANs. Devices would send a jam signal back to traffic sources when the device's receive buffers would start to overflow. The Ethernet CSMA/CD mechanism would sense the collision and stop transmitting for a pre-defined period.

Blocked Port

A port whose received traffic is stopped due to a condition in the system such as congestion or queue servicing problems.

Head-of-Line Blocking

A problem in which frames that are destined for other output ports must wait behind the head-of-queue frame. This condition arises when the frame at the head of an input queue cannot be dequeued due to congestion at its destination output port.

In Ethernet (RFC 2889), the method to determine head-of-line blocking is to examine the frame loss on the uncongested output port which is shown as Port 4 in the Head-of-Line Blocking Test diagram. Since Fibre Channel does not permit frame loss, the method for Ethernet does not apply. The method used in Fibre Channel switches is to measure the throughput at the uncongested output port (Flow C).

Intended Load

The number of frames per second that an external source attempts to transmit to a DUT/SUT for forwarding to a specified output interface or interfaces (RFC 2285). SmartFabric specifies this value as a percent of the wirespeed rate for the interface.

Configured Load

The load expressed in percent that represents the Smartfabric "Configured to Send" value. The "Configured to Send" value is the calculated number of frames that can be sent by each flow during the test interval. It is based upon the percent utilization specified for each port, the number of flows configured per physical port, frame size and port speed.

DUT

The Device Under Test.

FSPF

The Fabric Shortest Path First protocol that is the standard routing protocol for Fibre Channel Fabrics. Every switch in the fabric maintains a copy of the topology database and has knowledge of all ISLs in the fabric. The cost of a link is based on the link speed and the path with the lowest cost is used for frame delivery.

Ideal Load

The expected load per flow assuming fairness for each input port by the DUT.

ISL

Inter-Switch Link. A link directly connecting an E_Port of one switch to an E_Port of another switch.

Oversubscribing port

A port that may potentially request bandwidth in excess of what is available during an instant in time. This is not the same as a blocked port.

PortChannel A feature that allows multiple Fibre Channel ports to be aggregated into one logical link. It allows load-balancing across multiple ISLs at a lower level than FSPF load-balancing. Up to 16 ports, located anywhere in the system, may be included in a single PortChannel. PortChannel load balancing is based upon (S_ID and D_ID) or (S_ID, D_ID, and OX_ID).

RFC 2285. Benchmarking Terminology for LAN Switching Devices (February, 1998).

VSAN. Virtual SAN. The ability to create several isolated SAN topologies within the same physical infrastructure (analogous to VLANs). Each VSAN has its own set of fabric services.

Wirespeed. The maximum rate at which frames can be transmitted through an interface with the minimum interframe gap. The number of frames per second is dependent upon the frame size. The minimum interframe gap (IFG) specified in the ANSI standards is six transmission words, or more specifically, primitives. The primitive that most frequently occurs in the IFG is the IDLE fill word.

Zoning

The ability to partition the SAN into logical groupings of devices that can access other devices within the same controlled group. This feature differs from virtual SANs because all devices in a particular virtual SAN are still dependant on a common set of fabric services. These fabric services include name services, which provide the zoning functionality in each virtual SAN. Virtual SANs are used to partition fabrics and zoning is used to create subsets of devices within a virtual SAN.

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco.com Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy
Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2003 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0304R)